

Distributed Science Operations for JPL Planetary Missions

Richard D. Benson and Peter B. Kahn

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

N 94-23850

ABSTRACT

Advances in spacecraft, flight instruments, and ground systems provide an impetus and an opportunity for scientific investigation teams to take direct control of their instruments' operations and data collection while at the same time, provide a cost effective and flexible approach in support of increasingly complex science missions. Operations of science instruments have generally been integrated into planetary flight and ground systems at a very detailed level. That approach has been successful, but the cost of incorporating instrument expertise into the central mission operations system has been high.

This paper discusses an approach to simplify planetary science operations by distributing instrument computing and data management tasks from the central mission operations system to each Investigator's home center of observational expertise. Some early results of this operations concept will be presented based on the Mars Observer (MO) Project experience and Cassini Project plans.

1. INTRODUCTION

The Jet Propulsion Laboratory (JPL) has been conducting the mission operations of planetary exploration spacecraft for 30 years. The typical flow of work reflects processes that have been proven effective and reliable, various solutions to historic problems, and a mature organizational structure. Generally, spacecraft expertise has been developed at JPL while the science instruments originate at various investigator sites. Conduct of mission operations has often been supported at the most detailed level by the central JPL operations organization to maximize the density and quality of science data return and to ensure error-free operation of the spacecraft system.

The goal of a distributed mission information system architecture is to significantly improve the cost effectiveness and flexibility of mission operations. This is essentially a re-engineering task which unleashes and distributes computing power to enable processes to be accomplished more simply. The clearest cost saving areas are: 1) the cost avoidance of relocating science instrument and/or investigation teams to JPL for the mission duration (invariably years long and very disruptive of careers and families as well as expensive) and 2) cost avoidance of translating instrument and investigation knowledge from developer and user organizations to a separate, centralized operations organization. In a distributed system, the scientists are able to conduct most project-related science instrument operations from their home institutions via networked connections to a mission information system.

Improved flexibility comes from partitioning tasks and resources along investigation lines so that changes in investigation objectives and observing tactics can be done with comparatively minimal coordination with the central system. One benefit to a distributed operations system is that it provides for faster Science team response to changing instrument conditions. The data are provided to those with the most expertise and knowledge of the instrument without incurring significant time delays that would otherwise be imposed in a more centralized system.

MO is first JPL mission to use a distributed approach for science operations. This approach provides a unique "telescience" capability for the science instrument teams to remotely generate the observational uplink sequences and to provide near real-time analysis of the health and welfare of their instruments. Within the Space

Science community this is not a unique capability. The recent STS-52 flight of the Space Shuttle Columbia made use of telescience capabilities. Other missions, Hubble and UARS for example, have also utilized distributed science operations.

Cassini can utilize the experiences gained by MO and apply them to their unique needs to improve overall system effectiveness and flexibility. The Cassini mission will expand on this distributed approach capability for its operational requirements. Both missions involve significant on orbit time as opposed to one time flyby missions, so both can take advantage of the fact that observing opportunities are not once in a lifetime events but can be repeated on subsequent orbits.

Many operations tasks, especially system-level tasks, are more efficiently done centrally. However some subsystem functions are best distributed due to the 1) decreasing cost of computational power, putting more capability in the hands of the user, 2) the widespread use of communications infrastructures such as local and wide area networks and 3) financial constraints to keep staffing levels down and making better use of "smart" systems.

2. END TO END DESIGN CONCEPTS

Information system engineers at JPL try to think in terms of data processing in a true system sense. The central idea in "end to end" system engineering is to view processes and interfaces from an overall perspective to guide instrument, spacecraft, Deep Space Network (DSN), and Mission Control Center engineers toward designs that increase effectiveness and reduce costs by taking the end-user needs into account. Often this has the simple, but significant payoffs of eliminating redundant processes and promoting interface compatibility by making internal processing choices in favor of compatibility over expedience. An end to end view can also ease the development to operations transition of a project. Often that leverage is harder to apply because the phases are usually funded by different sources.

We see distributed science operations as a natural continuation of the end to end

approach. We discuss six areas where distributed science operations advance or are supported by the end to end approach. First is the delegation of function to whatever node is the best position to apply expertise directly. Second is processing, often applications software, that enable system and subsystem users to contribute to information processing tasks. Finally are database, standards, archiving, and security processes that supply an enabling infrastructure.

2.1 Decentralization of Science Operations

Cassini distributed science operations seeks to build on the experience of MO, which is the first JPL mission to use distributed instrument operations along with standard data interchange formats and multi-mission services to reduce project costs. There are differences in constraints between these two Projects, especially on the uplink processing. The distribution of tasks and trade-off analysis is customized for each mission. However, there are significant similarities in processing that lead to adoption of a distributed data system. Both missions take a hierarchical approach to operations. Both make significant use of multi-mission resources to build their uplink and downlink processes. Both distribute their operational processes and support to numerous teams, including remote science teams. MO has seven principal and six interdisciplinary investigations. Cassini has 12 principal and six interdisciplinary investigations. Most of these teams are not at JPL. The science teams further distribute the operational functions of the particular instrument to a variety of components within their own internal organization.

Some of the principles driving distributed operations are controversial. For example, the trend in many space science disciplines is to place observing decisions, including their technical and instrumental aspects, in the hands of observers who have operational expertise; this moves tasks away from historic, mostly centralized, venues.

Client-server architecture and widely available processing power have revolutionized data processing roles and made delegation of operational tasks to the

observer feasible and warranted. However, greater instrument/subsystem autonomy is not without some risk.

Packet command and telemetry standards allow for inheritance from mission to mission thereby decreasing development costs and time, but may involve more overhead (in terms of data volume) than custom formats.

The overall effect of the network and standards trends is to enable assignment of system functions to sites where they make the most sense. Certain functions that were done in centralized processes can now be delegated to the system's end users if it can be accomplished in a better, faster, cheaper process. Whether each function is done centrally -where common processes and system-level expertise can be applied, or distributed -where subsystem expertise can be directly applied, needs to be weighed. The emphasis of the distributed approach is on setting up a process that creates a satisfactory product in a minimum number of steps as opposed to a process that reshapes a product via rework into one that is ultimately satisfactory. Inheritance of design features from project to project decreases development and operational costs.

Coordinated efforts between the project mission operations teams, science teams, and multimission teams is essential for the design and operation of information data flow necessary to conducting planetary missions. MO had good success in hosting appropriate working groups. Cassini is just beginning to build such a working relation.

2.2 Processes

2.2.1 Sequence Design and Planning

The Science Investigation Teams are responsible for overall science operation (both scientific judgment and instrument operations aspects) because the requisite scientific, observational, and instrument expertise primarily reside there. Observing resources are allocated through a responsive, but systematic, process that resolves science conflicts at the science level based on science value judgments. In the end, all users are obligated to stay within their final allocations. The central

Ground System assigns reasonable margins to facilitate mission operations. (The term Ground System (GS) is used in the Cassini Project, and includes all applicable institutional and project-unique elements whose operation is centered at JPL. The GS includes the central functions in contrast to the distributed science functions.)

2.2.2 Command Generation

Instrument Commands - Detailed observation designs, both functional and performance aspects, are done by the Investigation Teams. Instrument-internal sequences/commands are generated, constraint-checked, and validated by the Investigation Team. The Ground System accepts these inputs as effective and reliable. The Ground System will not recheck these at a micro-level, in fact it generally can not. The GS may check them for compatibility with the system-level resources allocations and conventions depending on project-specific needs.

System Commands - The Ground System is responsible for mission success in an overall sense. This is a much broader role and includes providing:1) a physical environment that enables the instruments to successfully operate (e.g., trajectory, pointing, thermal, radiation, and supporting services); 2) central sequence planning, integration, uplink processing, and orchestration; includes tools for conflict and constraint checking; 3) central data collection on board, accounting, and transmission to the ground; 4) central data reduction, correlation, and distribution on the ground.

2.2.3 Instrument Health Assessment

Overall system-level analysis is the responsibility of the Ground System. This includes analysis of the spacecraft subsystem performance, the space-ground link, and the GS itself. This analysis can include instrument assessments up to their interactions with the overall system. It should be noted that the GS itself is a distributed system with various project-specific and multi-mission components.

Instrument subsystem-level performance analysis is the responsibility of the science

teams, although many capabilities of GS tools and teams can be applied. The cost and utility of the GS tools needs to be compared to the cost of using investigation-specific tools, such as the instrument support equipment retained and upgraded from the instrument development phase.

By providing the capability for the remote teams to monitor their instrument parameters in near real-time, engineering analysis and response times are faster.

2.2.4 Science Analysis

Science data analysis to interpret the observations, to decide future observing tactics, and for final science data record production are the responsibilities of the Investigator.

2.3 Centralized Project Databases

A core component of a distributed data system is a mechanism that allows for transfers of data files between various Project elements. One such component is a Project Database (PDB). The use of commercially available database systems tailored to the specific needs for a Flight Project is essential for the dissemination of data between operational components of the Project. Use of such a system allows for managed distribution and archival capabilities. A PDB is built and maintained as part of the multi-mission capability of JPL. The inheritance of the concept is improved from mission to mission. The use of specifically labeled "Data Objects" enables science and engineering teams to store and retrieve data products in a centralized PDB and eventually archive them in the Planetary Data System. Both raw telemetry and file types of data are stored.

The PDB design is divisible into two primary components; a relational catalog and a database for telemetry and file data. The database catalog is the primary mechanism for the Flight Project science and engineering teams to determine what data are available in the database. All data sets must be sufficiently self-described in the catalog to allow the PDB to uniquely identify the product. The Science and GS teams can query the Database for values

associated with the predefined attributes of a data product, as listed in the Catalog. Interactive or scripted selection of a displayed list of associated attributes for any chosen data object can be accomplished. The user can select and query for data in a systematic way allowing for routine operating procedures.

2.3.1 Science Team databases

The Science Investigation team have, in many cases, created central databases or distribution systems at their facilities in order to supply recently-acquired instrument data to their team members and Co-Investigators at remote locations from the instrument operations teams. Electronic access to these facilities is provided, through a contract with the Project, by the NASA Science Internet (NSI). This is the primary distribution mechanism for intra-team science products during the life of the mission. Alternately, the data residing in these databases are copied onto some form a media and mailed to the team members.

2.3.2 Science community/National databases

A central catalog of archived planetary data is available to the Planetary Science community through a distributed system of discipline database archives. Flight Project archival data products generated using a common metadata labeling standard are available to the general Planetary Science Community through these discipline nodes. It is important that the catalog information used to query for the individual products be provided to these central databases along with the archival products themselves. Scientists from around the world can request data from all the existing and past Planetary missions for research. The requester queries the central database to find where the specified data are located. The users then submits a request for that data with the discipline node. Access to this system enables scientists not associated with the specific Flight project to participate in the evaluation and study of the data collected by the mission.

2.4 NETWORK STANDARDS AND COMMERCIAL TECHNOLOGIES

The distributed nature of a ground data system for science operations is made possible through the use of standardized, commercial technologies, such as, local and wide area computer networks, the Internet infrastructure, off-the-shelf network communication working protocols and UNIX-based computer operating systems. Instead of relying on large centralized Project support teams, the Science Investigation Teams, located throughout the country, can utilize these technologies to perform their own analysis of the data without significant time delays.

Compact Disk- Read Only Memory (CD-ROM) as a data distribution media is another area of commercial technology being adopted by JPL Flight Projects such as Viking, Voyager, Magellan and Mars Observer. Use of the international CD-ROM ISO 9660 standards along with data format and self-description standards continues the evolution of common Planetary Science interchange methods for heterogeneous data sets.

2.5 ARCHIVING

Early definition and design of a Project's archive process is essential to ensuring the smooth flow of data to the general science community. Once again, the distribution of this task from a central organization to the instrument operational sites increases the knowledge base applied to the data and expedites the product deliveries to the general science community. With the advent of CD-ROM technologies and software, each science operations team can produce and deliver archival data products directly to their discipline archive node without the need for large, centralized Project archive and distribution teams. Cost efficiencies are gained as is the flexibility of the individual science team to tailor it's products specifically for the end user community. Use of common mapping formats and time tags is essential for cross-correlation and interdisciplinary studies. Early planning of the archive process in the design of the Project insures commonality of format and structure.

2.6 SECURITY ISSUES IN TODAY'S ENVIRONMENT

One of the major issues from a cost and schedule perspective when designing a distributed operations system is system security. The emphasis and institutional sensitivity placed on computer network violations has increased significantly over the years. A Project needs to balance the need for open interchanges of information between distributed nodes and users vs. the fear of intrusion that can force limited access to mission data even for validated users. This risk assessment or trade-off analysis and the associated cost impacts of open vs. closed systems is a challenge in the era of distributed operations.

The area of most sensitivity to Flight Projects is the security of the uplink system. This system in which an Instrument team sends sequence files or commands to an instrument in flight, poses the greatest risk to mission objectives. Obviously, an unauthorized intrusion into this system places at risk the whole mission. Substantial checks and balances by both the remote instruments teams and the institutional GS teams is required to ensure the validity of the sequencing system.

An interesting dichotomy is that the Project receives services from the NASA Science Internet (NSI) to connect all their remote science users. This access to the national infrastructure, called the Internet, both increases accessibility to the data along with promoting Interdisciplinary studies. At the same time this access puts the Project at some level of risk of network intrusions by non-authorized users.

A risk analysis needs to be performed that balances the needs of the free interchange of ideas and information against the need to prevent unauthorized intrusion and possible destruction of mission critical data. Quick expensive hardware or software solutions may not provide security and can inhibit necessary science and mission interactions.

Table 1. DIFFERENCES BETWEEN MARS OBSERVER AND CASSINI

Area	<u>Mars Observer</u>	<u>Cassini</u>
Mission Plan	Circular, sun sync, repetitive orbits	Orbital tour, 60 orbits, each somewhat unique
Specific Targeting	No, S/C Boresight always points at nadir	Yes, boresight is under constant, specific control
S/C Operating Modes	Simultaneous observations and earth communications	S/C either points remote sensors at target or antenna at earth
Orbit activity level	fairly constant	High at encounter, lower at apoapsis
S/C Power	Very robust, large margins	Logical restrictions (Operational Modes) needed to avoid overdraw
Cruise activity	Instrument checkout, maintenance, calibration	Limited maintenance during cruise, No instrument calibration until Jupiter (June 2002)
Instrument Command Verification	Validated commands are verified by GS, but not functionally simulated	All commands may be re-verified by GS via system level simulation or analysis (to be determined)

3. COMPARISON OF MO AND CASSINI

MO is currently operating enroute to Mars while Cassini is in development for an October 1997 launch and June 2004 arrival at Saturn. Both will spend years orbiting and observing their respective planets. Some examples of differences in the distributed operations approach of these missions are shown in Table 1. The most significant difference is in instrument command verification by the central GS. MO has a list of validated commands that have been proven in system testing or in-flight checkout to be safe, predictable, and effective. MO investigators are free to use these command at will with the GS ensuring that the header is correct. Cassini GS personnel currently see their spacecraft as more subject to instrument driven pointing and power conflicts. They are weighing whether every sequence should be validated through full simulated expansion of the instrument commands to ensure that the instrument sequences do not adversely affect the total sequence or spacecraft. This concern also motivates the current plan to leave the instruments turned off for the first five years of cruise, until adequate ground software can

be built to operate them. Clearly, this baseline plan will improve in favor of more free and early instrument operation as the Cassini teams become more comfortable with distributed operations.

4. CONCLUSIONS

The distributed data system approach for Planetary Science Operations provides more flexibility and allows tradeoffs to be made between the various components of the whole system in order to optimize the end-to-end data return and reduce project costs. Translations of technical expertise to another implementing or operating agency are reduced. Mission inheritance and adaptability are enhanced.

MO is having excellent success in their distributed operations. The Cassini Ground System Review Board finds, "...this is the correct approach for a mission of this type which will fly in the late 1990's." Our challenge is to build on the distributed approach for more responsive, lower cost operations.

Acknowledgement

This work was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.